




“... I envy eight year-old kids mastering the network. They have done it twenty years earlier than I. We had to invent it to do what they do...”

...“In the earliest days, this was a project I worked on with great passion because I wanted to solve the Defense Department's problem: it did not want proprietary networking and it didn't want to be confined to a single network technology”...

Vint G. Cerf Internet Pioneer.

The Internet history would not be understood without his contribution, without a doubt this summary reflects why Cerf receives the affectionate nickname of *Internet father*, although he insists on signing like "*Internet pioneer*". Vint, is a man who has devoted his life on solving complicated technological problems with simple ideas and brilliant solutions.



Interviewed¹ on December 8th, 2007 in McLean, Virginia (VA).

Born in June 23rd 1943 in New Haven, Connecticut in the Yale University Hospital.

My mother was living there with her father while my father was in the European theatre of World War II. He was in the navy and served as a gunnery officer on board a submarine chaser. Before the atomic bomb was dropped on Japan, he was preparing to be shipped out to the Pacific theatre where they had intended to invade Japan but then after the nuclear devices were dropped, the war was over and so he was out and we moved from New Haven to Los Angeles

in 1946.

So I actually grew up in the Los Angeles area. I was in the San Fernando Valley which is North of LA proper, in a little town called Van Nuys and eventually went to Van Nuys High School and that happened to be where Steve Crocker and Jon Postel and Karl Auerbach went to school. We weren't there at the same grade. I think Jon Postel was a year behind Steve and Steve was a year behind me. Steve and I were best friends in High School and we still are but neither of us knew Jon Postel then. We didn't actually meet Jon until we all became graduate students at UCLA.

1961: After I left high school I went to Stanford University and studied there from 1961 to 1965 where I majored in Mathematics and minored in German².


¹ Interview made from talks on June 25th 1999 (in San Jose, (CA); 20-May-2000 (Tarragona), 12-6-01 (Stockholm) and completed with electronic correspondence of days 23-6-01 (Tenerife) and 23-11-01 (Barcelona). Plus the recorded interview in McLean (Virginia) during the evening of December 8th 2007.

² I have been studying German since I was about 13 years old; my father invited a German tutor to come to the house once a week and I would learn German from him. I went to Stanford in Germany, a campus in an area around Stuttgart



1965: When I graduated from Stanford in June 1965 I went to work for IBM in what was called the Los Angeles Data Center. I was running a time sharing system for IBM called Quiktran. It was an interactive FORTRAN system so you could remotely use the mainframe IBM 7044 to write FORTRAN programs and execute them. We had people running stress analyses of buildings; we had one customer who was using the service to write algorithms for buying and selling stocks. There was just a wide range of applications, anything you could program in FORTRAN. I ran that operation center as a system engineer for 2 years until mid 1967 and then I took a leave of absence to return to school because I felt the need to get advanced training in computing.

1967: At UCLA I took a master and then PhD in computer science over 5 years period and my thesis advisor was Gerald Estrin. For Gerry I did a dissertation on multiprocess synchronization, multiple computers being run parallel but I also worked for Leonard Kleinrock as principal programmer for the Network Measurement Center and our task was either to inject traffic and to gather data from the ARPANET in order to validate the queuing model that Leonard and his students were developing for the packet switching networks. As a consequence, I was very quickly drawn to ARPANET activities in the late 1960's and ultimately worked together with Steve Crocker and Jon Postel on the first host to host protocol called NCP for Network Control Program. We implemented those and then demonstrated the system publicly in October 1972 in the basement of the Washington Hilton Hotel, in the International Conference on Computer Communications (ICCC). Bob Kahn organized the demonstration at the request of Larry Roberts who led the ARPANET program for ARPA and quite a few of the people that you have already interviewed were at that meeting including Louis Pouzin (Cyclades, France). At that meeting we formed a group called the International Network Working Group that ultimately become IFIP (International Federation on Information Processing) Working Group 6.1 which dealt with data communications. This group of people was quite interested in the notions of computer networking and packet switching in particular. At Steve Crocker's recommendation, I became chairman. Although I don't think any specific developments came out of this group as a working group, there were very useful discussions on network architecture and protocols.



where we stayed at a little town called **Beutelsbach**, which had about 3000 people, about 30 km away from Stuttgart and the house that we stayed in was a farm on top of a hill that was called **Landgut Burg** and there were about 70 students there. So Stanford brought in professors and classes were held there, sometimes in English but mostly in German. So we got to know the townspeople and traveled around in Europe during the six months that I was there, to Bonn, Prague, Rome, Stockholm and so on because we took advantage of weekends and it does not take very long to go from one part of Europe to another. I came back in December 1962 to continue my studies at Stanford.



- 1973: Bob Kahn and I presented to that group in September 1973, the basic design of the Internet in a document that we called INWG # 39. There was a number of working documents which came out of this group and a number of the people that you have interviewed³ were involved in this working group or documenting it acting as editor of the series of the notes. I don't know what happened to those notes today and it would be interesting to find out whether they are in the computing history museum or not.
- 1972: After the demonstration in October 1972, Bob Kahn left BBN (Bolt, Beranek and Newman) and went to ARPA and I left UCLA and went to Stanford, this time as a faculty member and joined both the computer science and the electrical engineering departments. That was where a lot of internet research began and was supported by DARPA and there were many other people you have interviewed who worked under Bob's sponsorship while I was teaching in Stanford.
- 1973: From 1973 through 1976 the TCP/IP protocols were brought up, the IP part wasn't actually split off until something like 1977 and it was at the urging of Danny Cohen, David Reed and Jon Postel that we made the separation of the IP layer protocol from the TCP layer protocol. It was primarily motivated by the desire to carry real time traffic that didn't necessarily have to all be delivered whereas TCP tries to deliver everything to the other end, in order and without duplicates. The real time traffic needed to be delivered quickly even though not all of it actually got there and it could even be delivered out of order if necessary. So the split of IP and TCP allowed us to put a base on packet transport including real time traffic. We used to test things like Voice over IP, which was just being looked at in the 1970's even though it only became visible and popular to the public⁴ in 1990, we were testing it in the 70's.
- 1976: After working at Stanford, ARPA asked if I would come to Washington DC to manage the internet research program for ARPA and so I agreed to do that and came in July or August 1976 to Washington and stayed there ever since, so it has been 31 years now. I undertook to manage the internet program, the packet radio program, the packet satellite program, and the security program, which I inherited from a man named Steve Walker, who was another ARPA program manager.

What is special about VOIP is that it's just another thing you can do on the Internet, whereas voice is the only thing - or nearly the only thing with the exception of the dial-up modem and fax - that you can do on the public switched telephone

³ A man named Alexander McKenzie may very well know what's happened to them and also I think Peter Sevcik would know because I think both of them at one time were editors of the series.

⁴ The first introduction to the public was by a company called Vocal Tech in Israel in 1996. Skype has been a very good example in public interest or demand in Voice over IP. My estimates now for the number of the internet user by the end of 2010 there would be possibly 3 billions users and the reason there would be so many of them is the use of mobile as the way to access the internet.



1982: I stayed at ARPA until late 1982 and then left the government to join MCI to develop what was eventually called *MCI-Mail*, which was a commercial email service that was launched on September 27 1983 and I stayed at that effort and continued to evolve it until June 1986 when Bob Kahn, who had left ARPA by that time decided to start a company called Corporation for National Research Initiatives (CNRI). That research organization was formed to focus on information infrastructure and Bob asked whether I would like to join him to do research and I agreed to do that so in June 86 I joined CNRI until October 1994 when I was invited to return to MCI in order to put them in the internet business. I agreed to do that and rejoined them as Senior Vice President for Data Architecture. Eventually I inherited the Internet Engineering activity at MCI and stayed with the company through a very painful period because the company was acquired by Bernard Ebbers and his WorldCom organization in 1998 and within four years Ebbers destroyed the company and put it into bankruptcy. June 2002 the company declared bankruptcy and I suppose I could have left along with a lot of other people but I thought it was ethical to stay and help put the company back together. I don't claim any great role in that but I figured I managed to do all I needed to do for the people who reported to me. I stayed until late 2005 at MCI and it was announced then that MCI would be acquired by Verizon and to be honest with you I felt that I had done a fair amount of work in the infrastructure of the internet for the last 11 years at that time and it would be sensible for me and interesting to move into application space again as I had with Bob Kahn in CNRI. It occurred to me that Google would probably be the most interesting place to go to for applications on the internet so I sent an email to Eric Schmidt who was CEO of Google and whom I have known for many years now and I just simply asked him if he wanted some help and he sent back a short note saying yes, so I agreed to join Google and I came on to the company in October 2005. I have been there since that time. We had a debate about what my title should be and when they asked what title I wanted I suggested Arch Duke; that seemed like a really great title but they reminded me that the previous Arch Duke was Ferdinand and he was assassinated in 1914 and it started World War I so that maybe being an Arch Duke wasn't good idea. They suggested considering what I have been doing for the last 30 plus years so I should become Google's chief internet evangelist, so that's my new title at Google and I spend an enormous amount of my time traveling in order to carry out that role.

www.ucla.edu
www.ibm.com
www.stanford.edu
www.arpa.mil
www.mci.net
www.google.com

Do you remember when you had your first contact with a computer?



It was in 1958, I don't recall what month of the year. It was in Santa Monica, California at a company called Systems Development Corporation, which was a spin out from RAND Corporation and RAND was doing classified research for US Air Force. Notably, Paul Baran at RAND Corporation wrote one of the first lengthy documents about what became packet switching; he didn't call it that at the time. RAND had been working on command and control for the national military and one of the ideas that came out of the RAND work was to put radars in the Northern part of Canada looking towards Russia and to transmit the radar information by land line all the way to Santa Monica, California where the data would be absorbed into a computing complex which was called the Semi Automated Ground Environment (SAGE) and that system was built with a tube based machine so the idea was the data would be transmitted back from the radars and would be displayed on large radar screens, that it would be processed by tube based computers. I remember being brought in to the center of the system; you literally walked in to the computer in order to use it; it was so big, it was rooms filled with vacuum tubes along the walls. And I remembered becoming completely mesmerized by the idea that you could have a computer dealing with things from such a long distance away. Not more than two years later Steve Crocker got permission to use some of the computers at UCLA so this would be somewhere around 1960 or so. We were either juniors or seniors at the Van Nuys High. Steve and I were best friends and he invited me to come with him to use UCLA Bendix G-15 computer. Shortly thereafter I went to Stanford as an undergraduate and of course immediately took all the computer science courses I could. We were using a Burroughs B 5000 machine which was upgraded to B5500, programming in ALGOL or what was called Burroughs ALGOL. So that was very exciting time for me, both at Stanford as undergraduate but also as high school student using computers at UCLA. I have these very vivid recollections of being exposed to computing and becoming completely infected with the excitement of using such machines.

What was your first contact/experience with Internet or ARPAnet?

First contact with ARPAnet comes with Leon Kleinrock at Network Measurement Center that would have been in 1969 or could have been as early as 1968. That was the year the RFQ was issued for the IMPS (Interface Message Processors - the packet switches of the ARPANET) and interestingly enough Steve Crocker and I bid on the ARPAnet project while we were still graduate students but we were working as consultants for a company called Jacobi Systems at the time. Jacobi bid on ARPAnet; it was a competitor of Bolt Beranek and Newman (BBN) and we ended up in the last few rounds, I think we were among the last four. I wrote a simulation on the packet switch net using the GPSS programming language but we didn't win and so by very good fortune, even though the company we were working for didn't win, Steve and I were still graduate students working in Len Kleinrock's group and so when BBN won, UCLA got the network measurement center contract and we got to work on that, so I got to work on the ARPAnet despite the fact that we didn't win on the original bid. Jacobi might have been the name of the founder of the company; they were based in Santa Monica.

On the internet side, it's already obvious because I began work the project in 1973 so that was clearly the first contact. When we started the project, we were



referring to things like internetworking and the idea was taking packet switched nets and connecting them together. But I remember seeing references to the shortened phrase internet as early as December 1974 or even earlier. The first specification of TCP protocol referred to the Internet TCP protocol and I think we already adopted in normal discourse the idea of these multiple network formed a thing called Internet, I always wrote that with capital I and as the net rolled out in 1983 and became available, people were saying “well, sometimes you could build pieces of network using Internet technology, which is not part of the public network.” So we said why we don’t call it with lower case internets, the networks that use TCP/IP protocols. It started out the upper case was used deliberately to refer to the ARPA project while the lower case were used to refer to networks that used the Internet technology but not connected to the DARPA Internet. Eventually when the Internet became public it was my habit to refer to the public Internet as the capital Internet and any other private networks that use the same technology as the lower case internets. But the term Internet was intended to reflect the idea that we had multiple networks that connect together, so the interconnections of the network and the framework in which that was done was the focus of attention of the research.

In your opinion, what are the key characteristics of Internet?

I think the most important characteristic is that it is *agnostic* when it comes to transport, it really doesn’t care whether the package is carried on optical fiber or on satellite channel or radio or landline or an Ethernet. It just doesn’t care and also doesn’t know what it is carrying. All it knows that it is carrying these internet packets but it doesn’t know what the bits in the packets mean. So the content in the applications of the internet are subject to interpretation by the computers at the end of the net, not by the network itself and I think that principle, the end to end principle is really important. What it has done is to support an explosion of applications that people had been able to write that are relatively insensitive to the basic internet architecture. They just ran over it so you don’t have to ask permission of an ISP to try out new ideas, you just do it and that’s a condition which creates huge incentives for new product development; you don’t have to get permission from anybody to try something out.

What do you consider the most important milestones in the development of the network?

- 1972-73 when Bob Kahn and I started talking about this
- 1974 December, first complete specification of TCP; it had bugs but Yogen Dalal, Carl Sunshine and I wrote that version. Very quickly thereafter implementation had begun and a series of redesigns had occurred.



- November 22nd 1977 we did a three network demonstration of the internet using the packet radio network in San Francisco bay area, packet satellite over the Atlantic and the wire line ARPAnet which at that time extended all the way into Norway and down to University College London.



So that demonstration showed that you can actually link all three networks together and have them perform successfully.

November 22nd 1977 at SRI (Menlo Park, CA) Courtesy of Don Nielson (SRI)

- 1983 is the major roll out of the internet in January 1st to all of the ARPA sponsored research community which requires the switch away from NCP to use the new TCP/IP protocols suite.
- 1989 first appearance of commercial services although I think the UUNet guys will argue that they had commercial internet services as early as 1987; they just didn't connect to NSFnet which I think is correct. The NSFnet first allowed commercial service on government backbone by 1989. Bob Kahn and I arranged to connect MCI Mail to the internet in 1988 and in June or July 1989 we actually implemented the MCI mail interconnection and we had it up and running in the middle of the year.
- The ARPANET was retired in 1990.
- The NSFnet which started in 1986 was retired in 1995 and so these milestones demonstrated that the internet has a life of its own. It has commercial services for those who are able to afford them and it didn't mean it was necessary to have research components in order to be operational.
- Tim Berners-Lee started the World Wide Web in 1988 or 89
- 1992 Marc Andreessen and Eric Bina from National Center for Supercomputer Applications wrote the Mosaic version of the WWW browser. Eric does not always get the kind of visibility he should have, a little bit like Robert Cailliau (CERN) in the case of the World Wide Web.
- 1994 Marc Andreessen (from NCSA) moved out to Palo Alto with Jim Clark (who started Silicon Graphics and Telemetry) and they started Netscape Communications. By this time I was back at MCI again and the guy I was reporting to at the time, Bob Harcharik and I flew out to Netscape communications to see if we could license their client technology for an electronic mall application that we were planning on doing at MCI and once again, MCI was about 10 years ahead of its time. We released MCI Mall with browser based interface and everything else and it does not get used very much because nobody is really aware of the web yet. But that was an important milestone because it was literally like an online shopping mall,



we had people put up web pages and you could go buy things and complete transactions online. It wasn't terribly successful and in some ways neither was the MCI mail project which started in 1982 and began operational in 1983 when there still weren't very many people accustomed to being online.

- 1996 the dot com boom was happening
- April 2000 the dot com burst when all of the investments and everything else, the internet start ups eventually fails because many of them didn't really have any business models at all.
- Then came the phoenix-like renaissance of the internet, during the period of 2006 to 2007 where there was a continuing demand for internet service quite independent of the burst of various internet application companies. The demand continued to grow at about 100 percent per year and today I would estimate it to be about 40 percent per year but that is still an enormous growth rate. I think that number is vastly inflated by mobiles that are becoming internet enabled. There are 3 billion mobiles on the market in 2007, 10 percent of them we believe are internet enabled, that's about 300 million and so if you add to that 500 million other devices on the internet, you got 800 million internet based devices so far. So that brings us more or less to the present where we are struggling with expansion of the internet, dealing with things like IPv6 and the addition of non Latin domain names into the internet by ICANN

How did you contribute to the development of the Internet?

In terms of my own contributions clearly the original design with Bob Kahn for TCP and the split to TCP/IP. I served on the IAB for many years and as its chair for a time. I founded the Internet Society along with Bob Kahn and others and served as its first president for 3 years. Eventually I joined the board of ICANN in 1999 and then served as its chairman from November 2000 to November 2007. The chair now is Peter Dengate Thrush who is an intellectual property barrister in New Zealand. He has been on the board for about two years now. So I think that I have made consistent contributions. Another one that is still underway is the operation of the internet across interplanetary space; that work has been going on since 1998 at the jet propulsion laboratory in Pasadena, California and I hope NASA and other international space agencies will adopt our Delay and Disruption Tolerant Networking protocols (DTN) for interplanetary communications, which are able to overcome some of the side effects of very long, uncertain delays in interplanetary communications. The delay could easily be hours when you get all the way to the outer planets. Mars is about 40 minutes round trip times at the speed of light at worst and about 6 minutes round trip time at best but when you get out to Saturn, Jupiter, Uranus and so on, you are talking about hours round trip time and a lot of uncertainties as well.

So I think that encompasses a lot of my contributions. At Google I am trying very hard to work on application side, trying very hard to help my colleagues at Google here to come up with new idea and to deal with new ideas in the internet applications.



I have actually been to the first of the internet governance forums, one in the Athens and most recently in Rio de Janeiro. I missed the most recent one in India. The people in the group and I contributed in that as well, thinking what internet governance actually means.

Who are some key people in the development of Internet, leaders or trendsetters?

- Related to the Internet:

Robert E Kahn (ARPANET and Internet); David Clark; Jon Postel; Robert Braden; Stephen T Kent; Dan Lynch (ARPAnet migration manager from NCP to TCP in January 1st 1983); Yogen Dalal, Carl Sunshine, Richard Karp, James Mathis, Ronald Crane, all them Stanford graduates who worked with the TCP in my Stanford laboratory; Dennis Jennings from NSF (National Science Foundation) for choosing TCP/IP for the NSFnet; Stephen Wolff (NSFnet), Hans-Werner Braun (MERIT - NSFnet); Gerard LeLann (worked at Stanford on the TCP design even he was from IRIA -France-); John Shoch and Bob Metcalfe who worked at Xerox PARC and came to my Stanford seminars about TCP); Peter Kirstein and his students from UCL (Univ. College of London); Ray Tomlinson and Bill Plummer from BBN; Noel Chiappa (MIT) worked in routers; Virginia Strazisar (she designed the first gateway - in BBN)... this is a very LONG list... especially from the most recent people from 1983. Larry Landweber and David Farber (CSnet); David Crocker, John Vittal, Ray Tomlinson - for designing one of the first email services for the ARPANET and Internet.

- From the ARPANET times:

Stephen D Crocker (for his pioneer work on TCP's predecessor: the NCP which worked during the origins of the ARPANET until 1983); Larry Roberts (ARPANET), Len Kleinrock and Howard Frank (for the ARPAnet topology definition); Frank Heart and all the BBN team for building and designing the first IMP and network; Donald W. Davies, from the National Physical Laboratory, London, UK (as one inventor of packet switching) ; Paul Baran (RAND, USA) another inventor of packet switching.

A couple of anecdotal situations

One was with the packet radio network which was being implemented and tested by SRI international. We did a lot of demonstration of the packet radio system interconnecting with computers in the ARPAnet. Whenever I was present for those demonstrations, I was able to tell whether the packet radios were working or not because my hearing aid picked up the signal. It was transmitting in 1710 to 1850 MHZ band and my hearing aids were sensitive to those frequencies so I would hear these crackling paper sounds whenever packet was transmitted and as soon as the sound stopped, I knew that the radio had died and I would kick the SRI guy under the table again to restart the radio, I knew that it just stopped. This was typically sitting in a conference room with a packet radio inside all those. I just need to be within a few feet of the radio to detect the packet being transmitted so that was kind of amusing.



What do you think about the future of Internet?

As far as the future of the internet goes, I am convinced that it will continue to expand, there would be eventually more than 5 billion users, 6 billion maybe, I think it will probably take until 2015 to reach 5 billion, that's only 8 years from now but if I am right about the 3 billion in 2010, it gives us another 5 years to get the other half on and a lot of them because the mobiles are increasingly Internet-enabled so I don't think that's entirely impossible.

I also think that we are going to see all of the media show up on the network; radio, television, print and so on; it's already happening so it's predictable. I think people are going to use the internet to control their appliances so all household appliances and appliances in the cars and in the office and the things you carry around with you will be internet enabled. I think we are going to have to struggle to figure out how to index and interpret digital content on the network. I am worried about losing our ability to interpret the bits. If you have files for the bits but you don't have the program that knows how to interpret them then all the data and images you have won't be meaningful anymore, they will just be bits, which would be a terrible thing. So finding a way to make sure we can continue to interpret bits that are held on the internet is important and it probably means not just trying to upgrade copies by copying them with new applications but also to hold on to older application copies that would otherwise have expired. Somebody says "I am not going to support this application anymore" we, the community of internet users would like to have access to that software so we can make sure it still works because otherwise our files are no longer accessible. That raises intellectual property questions among other things, but I think we have to work that out because otherwise we will end up in what we call bit rot.

Do you see any technological trends?

The ones that are very clear right now are increased mobility, increased bandwidth at the edges of the internet, increased symmetry in access to network services; instead of having the asymmetric broadband we have now the US some bandwidth in symmetric form. I think there would be more mobile and I am hoping there will be a lot of interplanetary interaction with the internet over time.



Vint Cerf with Andreu Veà during the interview on December the 8th at McLean (VA)

ADDITIONAL READING

PAPERS & BOOKS MENTIONED / RECOMMENDED

Cerf & Kahn. "A Protocol for Packet Network Intercommunication". IEEE Transactions on communications Vol. com-22, no.5, May 1974.